A SAT model to mine flexible sequences in transactional datasets

Remi Coletta
LIRMM
University Montpellier 2
Montpellier, France
coletta@lirmm.fr

Benjamin Negrevergne
Inria
Centre Rennes, France
benjamin.negrevergne@inria.fr

arXiv:1604.00300v1 [cs.AI] 1 Apr 2016

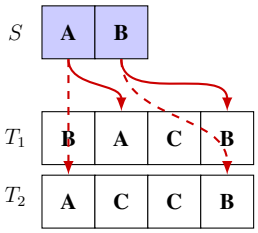

Fig. 1. Sequence $S = [A, B]$ is embedded in transaction T_1 (solid arrows) and transaction T_2 (dashed arrows).

***Abstract*—Traditional pattern mining algorithms generally suffer from a lack of flexibility. In this paper, we propose a SAT formulation of the problem to successfully mine frequent flexible sequences occurring in transactional datasets. Our SAT-based approach can easily be extended with extra constraints to address a broad range of pattern mining applications. To demonstrate this claim, we formulate and add several constraints, such as gap and span constraints, to our model in order to extract more specific patterns. We also use interactive solving to perform important derived tasks, such as closed pattern mining or maximal pattern mining. Finally, we prove the practical feasibility of our SAT model by running experiments on two real datasets.**

I. Introduction

Pattern mining, now both an important research topic and a broadly-used analysis tool, aims at extracting recurring patterns in large datasets. Since seminal work by Agrawal and Srikant on frequent itemset mining [1], a large number of new pattern mining tasks have been proposed to extract more *structured* patterns, such as sequence, tree or graph patterns. While structured patterns are usually more informative, they are also harder to extract.

The sequence mining problem is a variation of the itemset mining problem, in which both transactions and patterns are ordered. Given a set of ordered transactions, the goal is to find all the sequences that are embedded in more than a given number of transactions. The complexity of the task stems from the fact that we allow sequences to be embedded in the transactions with variable-length gaps. See Figure 1 for an illustration.

This problem has been given multiple names in the literature: depending on the application, it has been called *embedded subsequent mining*, *flexible motifs mining*, or *serial episode mining*. Nevertheless, it particularly affects biomedical applications because of the ubiquity of sequential (DNA or protein) datasets in this domain (e.g.,[2]). It also impacts text analysis, where sentences can be seen as sequences of words (e.g.,[3]).

Sequence mining, as well as other pattern mining tasks, was first developped by writing algorithms in classical imperative languages such as C or C++. In order to handle large datasets, most of these algorithms were optimized for specific mining applications. As a result, they are unable to capture small variations. This is a major drawback for users, whose needs may substantially vary from one application to the other. For example, text analysts generally have grammar constraints, whereas biologists usually require very specific, per-character constraints.

To solve this problem, Guns et al. in [4] have recently proposed using CP modeling techniques and CP solvers to address various itemset mining tasks. In contrast with imperative programs, CP models can easily be extended to efficiently perform a variety of specific mining tasks, while remaining more or less performant on the basic mining task. However, to date, most of the published works have focused on modeling and solving set-based structures, which are natively supported by solvers or easy to encode as boolean vectors. Nevertheless, the question of the usability of ... needs to be further examined. Whether the same solvers can be used to mine more structured patterns, such as sequences, trees or graphs using generic solvers, remains an open question.

a) Contribution of this paper: In this paper, we demonstrate that the sequence mining problem can be expressed as a purely boolean SAT formula, and solved using off-the-shelf SAT solvers.

To extend the basic problem and address more specific needs, we formulate several constraints commonly used in sequence mining. We show how these additional constraints can be easily incorporated into our SAT formula and without previous SAT knowledge. This is a critical benefit for users, who can now arbitrarily combine sets of constraints without concern for the soundness of their results. In addition, our experiments on real datasets establish that, contrary to specialized algorithms, SAT solvers take advantage of these

new constraints to further prune the search space and increase efficiency.

Lastly, we demonstrate that, in addition to enumerating *all* frequent sequences, our model also allows us to perform important derived mining tasks such as the *closed pattern* and the *maximal pattern* mining tasks using interactive solving. These two tasks aim to reduce redundancy among result sets and have been extensively studied in the pattern mining community.

The remainder of this paper is organized as follows: we formally state the problem in Section II. We then present our SAT model to mine flexible sequences in Section III and formulate several user constraints in Section IV. We propose our solution to mine maximal and closed patterns in Section V, and present our experiments in Section VI. Finally, we discuss related work in Section VII and offer our conclusions in Section VIII.

II. Frequent sequence mining: formal problem statement

Given:

1) a vocabulary $V : \{v_0, \ldots, v_{|V|}\}$
2) a set of transactions $\mathcal{T} = \{T_1, \ldots, T_{|\mathcal{T}|}\}$ where each transaction T_i is a sequence of characters in V; we use the projection notation $T_i[j]$ to denote the j^{th} character of the transaction T_i; and
3) a minimum frequency threshold $minsup$,

the sequence mining problem consists in enumerating all the sequences that are *embedded* in at least $minsup$ transactions.

In the standard sequence mining problem, we say that a sequence S of n characters is embedded in a transaction T if there exist $j_1, \ldots, j_n$ integers such that:

$$j_1 < \ldots < j_n, \forall i \in 1, \ldots, n \ s.t. \ S[i] = T[j_i]$$

The list of positions $e = (j_1, \ldots, j_n)$ is called an *embedding*. To express that e is an embedding of S in T, we note $S \sqsubseteq_e T$. We also note $S \sqsubseteq T$ if there exists at least one embedding e such that $S \sqsubseteq_e T$.

We now define the *coverage* of a sequence S as the set of transactions for which there exists an embedding of S:

$$cover(S, \mathcal{T}) = \{T_i \in \mathcal{T} : S \sqsubseteq T_i\}$$

Using these notations, the problem of mining all frequent sequences in a set of transactions can be formalized as follows: Given a vocabulary V, a set of transactions $\mathcal{T}$ and a minimum threshold $minsup$, enumerate all sequences S such that:

$$|cover(S, \mathcal{T})| \geq minsup$$

III. SAT encoding of the base model

In this section, we present the variables and constraints of our SAT model. We then show how to slightly modify the SAT solver in order to generate all the frequent sequences.

A. Variables

In frequent sequence mining, patterns (i.e., frequent sequences) can have varied lengths, whereas constraint-based solvers usually require a fixed number of variables. Therefore, one must first compute an upper bound for the size of the patterns. Obviously, a pattern cannot be longer than the longest transaction, but we may compute a tighter bound: to be frequent, a pattern has to be covered by at least $minsup$ transactions. Therefore, the maximal size K of a pattern is the size of the $minsup$ longest transactions in $\mathcal{T}$. We then define every possible pattern S as a sequence of exactly K characters and introduce an extra character ϵ in the vocabulary V to encode shorter patterns. $\mathcal{V} = V \cup \{\epsilon\}$. For example, if $K = 4$, the pattern ab will be represented as $S = ab\epsilon\epsilon$.

We can now build our formulation around 3 sets of variables: the m, c and t variables. See Figure 2.

m variables: For each position $k \leq K$ in the pattern S and for each character v of the vocabulary $\mathcal{V}$, we introduce a literal $m_{k,v}$, which is true if and only if the k^{th} character of the pattern is equal to v:

$$m_{k,v} = \top \Leftrightarrow S[k] = v$$

c variables: For each transaction T_i, we add a literal c_i, which is true only if the transaction T_i is covered by the pattern:

$$c_i = \top \Rightarrow S \sqsubseteq T_i$$

t variables: For each transaction T_i, for each position j in T_i and for each position $k \leq j$ in S, we add a literal $t_{i,j,k}$, which is true if the j^{th} character of the transaction T_i is used as support for the k^{th} character of S. If $c_i = \top$, there exists, by definition an embedding e such that $S \sqsubseteq_e T_i$. If $S \sqsubseteq_e T_i$ then $t_{i,j,k} = \top$ only if $e[k] = j$.

$$t_{i,j,k} = \top \ \wedge \ S \sqsubseteq_e T_i \Rightarrow e[k] = j$$

Note that we use an implication rather than an equivalence because there may be multiple matches for the character $S[k]$ in the transaction T_i. In Section **??**, we explain how extra user constraints prevent us from selecting matches *a priori*.

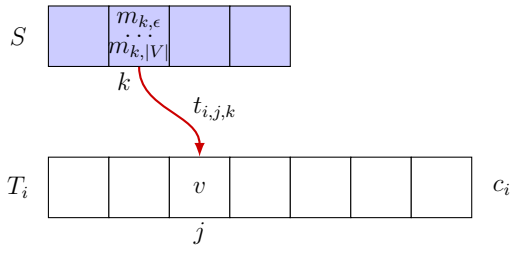

Fig. 2. Variables of the model

B. Constraints

*b) **Well-formed pattern**:* The two following constraints encode the fact that each position k in the pattern can only be assigned exactly one character from the vocabulary. Constraint (1) imposes that *at least one* character is assigned per position

since constraint (2) imposes that *at the most one* character is assigned per position:

$$\forall k \in 1..K, (m_{k,\epsilon} \vee m_{k,1} \vee \ldots \vee m_{k,|V|}) \quad (1)$$

$$\forall k \in 1..K, \forall v \in 1..V\text{-}1, \forall v' \in v\text{+}1..V, \quad (\overline{m_{k,v}} \vee \overline{m_{k,v'}}) \quad (2)$$

To avoid pattern symmetries, we discard every solution in which at least one ϵ character does not appear at the end of the sequence. In other words, we forbid solutions in which ϵ is followed by any character $c \in V \setminus \{\epsilon\}$.

$$\forall k \in 1..K\text{-}1, (\overline{m_{k,\epsilon}} \vee m_{k+1,\epsilon}) \quad (3)$$

*c) **Support compatibility**:* This constraint imposes that positions in the pattern can only be mapped to the positions in the transaction if the characters are equal.

$$\forall i \in 1..|\mathcal{T}|, \forall j \in 1..|T_i|, \forall k \in 1..i \\ (\overline{t_{i,j,k}} \vee m_{k,v}), \text{ where } v = T_i[j] \quad (4)$$

*d) **Coverage constraint**:* For each transaction T_i, $S \sqsubseteq_e T_i$ is covered by the pattern ($c_i = \top$) only if, for each position k of the pattern, either $S[k] = \epsilon$ ($m_{k,\epsilon} = \top$) or there exists a position j such that $e[k] = j$.

$$\forall i \in 1..n, \forall k \in 1..K\text{-}1, (\overline{c_i} \vee m_{k,\epsilon} \vee t_{i,k,k} \vee \ldots \vee t_{i,|T_i|,k}) \quad (5)$$

To avoid symmetries in embeddings, each character in the pattern can match only once in the transaction.

$$\forall i \in 1..n, \forall k \in 1..K\text{-}1, \\ \forall j \in 1..|T_i|\text{-}1, \forall j' \in j\text{+}1..|T_i|, (\overline{t_{i,j,k}} \vee \overline{t_{i,j',k}}) \quad (6)$$

*e) **Order preservation constraint**:* For each transaction T_i, for each position k of the pattern, position j is a valid support for $S[k]$ ($t_{i,j,k} = \top$) only if it exists $k-1 \leq j' < j$ s.t. position j' is a valid support for $S[k-1]$ ($t_{i,j',k-1} = \top$):

$$\forall i \in 1..n, \forall j \in 1..|T_i|, \forall k \in 1..j, \\ (\overline{t_{i,j,k}} \vee t_{i,k-1,k-1} \vee t_{i,k,k-1} \vee \ldots \vee t_{i,j-1,k-1}) \quad (7)$$

*f) **Cardinality constraint**:* The last constraint of the sequence mining problem is the minimum frequency threshold $minsup$. To ensure that pattern S is frequent enough, we need to encode the requirement that at least $minsup$ transactions $T_i \in \mathcal{T}$ are covered ($c_i = \top$).

$$atLeast(minsup, [c_1, \ldots, c_{|\mathcal{T}|}]) \quad (8)$$

This constraint appears in numerous industrial applications, and several decompositions in CNF have been proposed in the literature. Bailleux and Boufkhad [5] proposed a network-based decomposition using $\mathcal{O}(|\mathcal{T}|.log^2(|\mathcal{T}|))$ extra literals. This number of extra literals was later reduced to $\mathcal{O}(|\mathcal{T}|.log^2(minsup))$ in [6], and proved optimal. More recently, Abio and Stuckey [7] have proposed a technique to improve the efficiency of the cardinality constraint, by taking into account two tricky internal mechanisms of a CDCL SAT solver. Specifically, the authors suggested using the explanation of conflicts and the activity of variables after the restarts. However, in some variants of our problems (presented in Section IV), the cardinality between $minsup$ and $|\mathcal{T}|$ must be preserved. And because of the logarithmic formulation, there is no simple way to express sub-cardinalities in these approaches. We therefore used the formulation initially proposed by Coquery et al. in [8]. Although this decomposition involves more variables $\mathcal{O}(|\mathcal{T}|^2)$, it allows the simple expression of cardinality. Indeed, when building the decomposition for $|\mathcal{T}|$, we obtain a literal $card_t$ per value $t \in minsup..|\mathcal{T}|$, stating that if $card_t$ is true, there exist at least t literals among $c_1, \ldots, c_{|\mathcal{T}|}$ which are true.

C. Solving

We solve our SAT formula using the Glucose CDCL solver [9]. Truc et Machin have shown in [10] that branching over decision variables confuses the clause-learning system and ultimately reduces the solving performances. Therefore, we do not provide any hint for the search strategy and use the solver as a black box.

Each model of our SAT formula (i.e., each solution) is a frequent sequence of the input dataset. As a result, outputting all the frequent sequences can be done by enumerating all the models that satisfy the formula.

Algorithm 1: COMPUTE ALL FREQUENT PATTERNS

```
1 while F is solvable do
2     model ← nextModel(F)
3     F ← F ∧ (⋁ m_{k,v} s.t. model[m_{k,v}] = ⊤ ∧ v ≠ ϵ)
```

However, more specifically, only the literals $m_{k,v}$ of the model correspond to a frequent sequence; indeed, literals c_i represent the set of covered transactions, and literals $t_{i,j,k}$ an embedding of this sequence in each covered transaction. Since each covered transaction may have several embeddings, the number of solutions for our SAT formula may be exponentially larger than the number of frequent sequences.

To eliminate this potential issue, we propose Algorithm 1 as a simple way to enumerate all the models: we first solve the formula and find a model (line 2); we then restrict the model to literals encoding the frequent sequence, and subsequently post the negation of the restricted model as a new clause (line 3). The process is repeated until no additional model can be found (line 1).

IV. EXTENDING THE MODEL WITH USER CONSTRAINTS

Enumerating all frequent sequences may yield an overwhelming number of sequences. (See second column in Table II.) Some of them may be deemed redundant and/or irrelevant by the user. In this section, we formulate various user constraints that can be added to the base model in order to extract fewer and more specific sequences. The following is not an exhaustive list of useful user constraints, but rather a series of examples that demonstrate the flexibility of the SAT formulation.

A. Constraints on the way a pattern is embedded

*g) **Sequence mining with maximum gap constraint**:* With the maximum gap constraint, transactions are covered only if the characters of the sequence occur in the transaction with a gap of at most γ characters. We define the $max\text{-}gap$ predicate over an embedding $e = (k_1, \ldots, k_n)$ as follows:

$$max\text{-}gap_\gamma(e) \equiv \forall i \in 1 \ldots n\text{-}1, (k_{i+1} - k_i) \leq \gamma \quad (9)$$

To expand the standard sequence mining problem with this constraint, we only have to append the following constraint to the formula.

This constraint imposes that the distance between indices j and j' of $t_{i,j,k}$ and $t_{i,j',k}$ is smaller than γ:

$$(\overline{t_{i,j,k}} \vee t_{i,max(k-1,j-\gamma),k-1} \vee \ldots \vee t_{i,j-1,k-1}) \quad (10)$$

h) Remark on the gap constraint: Due to its usefulness, the gap constraint has already been implemented in LCMseq algorithms [11]. Unfortunately, though correct, this implementation is not complete. The gap constraint seems to be checked only on *the first embedding* of the sequence in each transaction. We illustrate this shortcoming with the following dataset, using $minsup = 2$ and $max\text{-}gap = 2$:

T_1 : $ACCBAB$
T_2 : AB

In this dataset, LCMseq is unable to find AB as a frequent sequence because the first embedding of AB in T_1 (A_1B_6) does *not* satisfy the gap constraint.

Indeed, counting all the transactions that contain *at least one* embedding satisfying the gap constraint (as formulated in Equation 9) requires a more complex algorithm.This short example points out the difficulty to integrate extra constraints in a traditionnal mining algorithms.

In contrast, our approach, based on an SAT solver, tackles this issue in a sound and complete way.

*i) **Sequence mining with dependent gap constraint**:* With the generalized gap constraint, transactions are covered only if the two consecutive sequence characters occur in the transaction separated by a gap depending on the position in the sequence and/or the character preceding the gap. More formally, given a function $gap : (K \times V) \rightarrow [1..l]$, we define the $gen\text{-}gap$ predicate over an embedding $e = (k_1, \ldots, k_n)$ as follows:

$$gen\text{-}gap_{gap}(e) \equiv \forall i \in 1 \ldots n\text{-}1, (k_{i+1} - k_i) \leq gap(i, S[i])$$

This constraint may be enforced by the following clause:

$$(\overline{t_{i,j,k}} \vee (\overline{m_{k,v}} \vee t_{i,max(k-1,j-gap(k,v)),k-1} \vee \ldots \vee t_{i,j-1,k-1})$$

*j) **Sequence mining with maximum span constraint**:* With the maximum span constraint, transactions are covered only if the first and last characters of the sequence occur in the transaction separated by at most γ characters. More formally, we define the $max\text{-}span$ predicate over an embedding $e = (k_1, \ldots, k_n)$ as follows:

$$max\text{-}span_\gamma(e) \equiv \forall i \in 1 \ldots n\text{-}1, (k_n - k_1) \leq \gamma$$

To encode this constraint, we define two additional literals per character in each transaction: $f_{i,j}$ points to the position of the support for the first character of S, and $u_{i,j}$ is true for each position j in the transaction T_i between the support for the first and the last character of S. $f_{i,j}$ and $u_{i,j}$ are subjected to the following constraints:

$$\begin{cases} \forall i \in 1..|\mathcal{T}|, \forall j \in 1..|T_i|, (\overline{t_{i,j,0}} \vee f_{i,j}) \wedge (t_{i,j,0} \vee \overline{f_{i,j}}) & (i) \\ \forall i \in 1..|\mathcal{T}|, \forall j \in 1..|T_i|, \forall k \in 1..j (\overline{t_{i,j,k}} \vee f_{i,j}) & (ii) \\ \forall i \in 1..|\mathcal{T}|, \forall j \in 1..|T_i|, (\overline{u_{i,j}} \vee f_{i,j} \vee u_{i,j-1} & (iii) \\ \forall i \in 1..|\mathcal{T}|, \forall j \in 1..|T_i| - span, (\overline{c_i} \vee \overline{u_{i,j}} \vee \overline{u_{i,j+span}}) & (iv) \end{cases}$$

Clauses (i) ensure that if $Ti[j]$ is used as the support for the first position of sequence ($t_{i,j,0} = \top$), then is marked as the first support $f_{i,j} = \top$; clauses (ii) mark any position k in the sequence ($\forall t_{i,j,k} = \top$, we have $u_{i,j} = \top$); clauses (iii) enforces positions i', which are not used as a support themselves but located between the first support $f_{i,j}$ and support and another support, to also satisfy $u_{i',j}$. Finally, (iv) enforces a gap of at most $span$ characters between the support of the first and the last character of S.

B. Constraints over the pattern itself

In many application fields, interesting sequences can be distinguished from irrelevant ones by specifying syntactic constraints on the pattern itself. For example, the authors of [12] designed an algorithm to mine frequent sequences that satisfy regular expressions specified by the user. The resulting algorithm is applied to biological datasets.

We now explain how to address this type of constraint in our SAT formulation. In [13], Pesant proposed the *regular* global constraint for CP systems. Indeed, this constraint is satisfied if a sequence of values belongs to a given regular language. More recently, in [14], [15], the authors extended this work to the context-free languages with the introduction of *grammar* global constraints. Furthermore, Quimper and Walsh [16] and by Katsirelos et al. [17] have shown how to decompose regular and grammar constraints into SAT formulae. Based on their work, we can easily incorportate expression and grammar constraints into our SAT formulation.

V. Maximal and closed pattern mining

In our SAT formulation and in other CP/SAT formulations of pattern mining problems (e.g., [4]), each solution represents a pattern. Usually, in classical SAT or CP applications, satisfiability (i.e., finding one solution only) is the main concern. Conversely, pattern mining users are not interested in finding a single solution, but multiples ones. In Section III-C, we showed how to enumerate all the solutions with minor modifications in the SAT solver.

However, the set of all patterns is usually highly redundant. To overcome this issue, condensed representations were introduced. For example, closed pattern mining [18] aims at extracting a subset of patterns from which it is possible to derive all other patterns. Machin et Truc [19] later showed that expressing closedness as a constraint can lead to ambiguous problem definition. In this section, we address closed and maximal pattern mining [20] by using interactive SAT solving.

*k) **Closed sequence mining**:* This was first introduced by Pasquier et al. for itemset mining ([18]).

Given the frequent sequence $S = v_1 v_2 \dots v_n$, we say that S is closed if there is no character v' such that $S' = v_1 v_2 \dots v_n v'$ has the same coverage.

$$closed(S, \mathcal{T}) \equiv$$

$$\nexists v' \in V \text{ s.t. } cover(S, \mathcal{T}) = cover(Sv', \mathcal{T})$$

In our base formulation, constraints 4 and 7 prevent the validity of supports $t_{i,j,k}$ due to the compatibility and the order of supports. Then, constraint 5 prevents the c_i from being satisfied, due to the validity of supports $t_{i,j,k}$ (c_i only appears as a negative literal). The only constraint that mandates that c_i be satisfied is the frequency constraint 8. As a result, given a pattern S, if $|cover(S, \mathcal{T})| > minsup$, there may exist transactions T_i such that $S \sqsubseteq T_i$ but $c_i = \bot$. This is captured by the definition of c_i given in Section III. Unfortunately, the definition of $closed(S)$ requires the maximization of $|cover(S, \mathcal{T})|$, which would only be possible if we had $c_i = \top \Leftrightarrow S \sqsubseteq T_i$ instead of $c_i = \top \Rightarrow S \sqsubseteq T_i$. However, coverage, order preservation and other user-defined constraints (*see* Sections IV-B and IV-A) only impose restrictions on supports. Having $c_i = \top \Leftrightarrow S \sqsubseteq T_i$ would require that we combine the negation of each of these constraints, then build a DNF; this would result in an exponential number of clauses.

To avoid this constraint negation requirement, we now propose to use the incremental SAT solving paradigm [21]. Algorithm 2 performs several calls to the SAT solver of formulae sharing the same set of clauses, but with a different set of assumptions.

The main advantage of SAT solving under assumptions is that the clauses learned during a solving procedure under a given set of assumptions are kept for later solving. The learned clauses are removed only if they involve a literal that occurs in an assumption that is also removed.

Since nested loops (lines 1 and 2) ensure that for any pair of patterns S_1 and S_2 found, with S_1 found before S_2, we have either $covers(S_1, \mathcal{T}) \subset covers(S_2, \mathcal{T})$ or $covers(S_1, \mathcal{T}) = covers(S_2, \mathcal{T})$ and $|S_1| \geq |S_2|$. For each pattern S found, line 4 adds no-goods corresponding to each subsequence $S' \sqsubseteq S$ to avoid generating S' in a future resolution.

Algorithm 2: COMPUTE ALL CLOSED PATTERNS

```
   Stack : A ← {card_{minsup+1}, ..., card_{|T|}}
1  do
   |  Stack : A' ← {¬m_{1,ε}, ..., ¬m_{K,ε}};
2  |  do
3  |  |  while F is solvable under assumption(A ∧ A') do
   |  |  |  model ← nextModel(F);
4  |  |  |  foreach pattern ∈ P(m_{k,v} s.t. model[m_{k,v}] =
   |  |  |  ⊤ and v ≠ ε}) do
   |  |  |  |  F ← F ∧ (⋁_{m_{k,v} ∈ pattern} ¬m_{k,v} ∨ ¬m_{|pattern|,ε})
   |  |  A'.pop()
   |  while A' ≠ ∅;
   |  A.pop()
   while A ≠ ∅;
```

Dataset	# transactions	voc size	max length	avg length
JMLR	788	3847	231	96.997
JMLR-500	500	3208	231	95.754
Gazelle	29369	1424	651	2.981
Gazelle-500	500	1094	651	15.052

TABLE I
CHARACHTERISTICS OF *Gazelle* AND *JMLR*.

*l) **Maximal sequence mining**:* Mining maximal sequences aims to extract only the largest sequences that remain frequent. This further reduces the number of generated sequences, but the set of all frequent sequences cannot be derived from the set of maximal sequences.

$$maximal(S, \mathcal{T}) \equiv$$

$$\nexists v' \in V \text{ s.t. } |cover(S.v', \mathcal{T})| \geq minsup$$

We address the problem of finding maximal sequences with a minor modification in Algorithm 2 by only removing the external loop; indeed, we are not concerned with maximizing the number of supported transactions as we were for closed patterns.

VI. EXPERIMENTS

After presenting our experimental setting in Section VI-A, we now compare our approach to specialized algorithms designed to mine closed frequent sequences. Subsequently, we demonstrate, that unlike specialized algorithms, our SAT-based approach is able to take advantage of extra constraints to further prune the search space.

A. Experimental setting

We run our experiments on two real datasets called *Gazelle* and *JMLR*. *Gazelle* is a dataset of web clicks data, in which each transaction represents the sequence of page views by a single user. *JMLR* consists in a set of paper abstracts from the *Journal of Machine Learning Research* in which each transaction is the abstract of a publication. Both datasets were previously used to evaluate sequence mining algorithms, for example in [22] or in [3]. Since our approach is currently unable to tackle the entire datasets, we derive smaller datasets by taking the first n transactions. For example, the *JMLR-500* dataset consists of the 500 first transactions of *JMLR*. This approach is preferable to synthetic datasets, which usually exhibit very different properties than real datasets. As shown in Table I, the properties of *JMLR-500* and *Gazelle-500* are similar to the ones of *JMLR* and *Gazelle* respectively.

When the frequency threshold is lower, the number of sequences to enumerate is larger. To solve our SAT formulation, we use the Glucose[9] SAT solver modified as explained in Algorithm 2. We only mine for closed sequences, since non-closed sequences can be derived from closed ones. We run all our experiments on a Linux host with an Intel Core i7-2600 CPU at 3.40GHz and 16 GiB of memory.

Dataset@support	# clo. patterns	LCMseq	Bide
Gazelle-500@10%	1	0.01s	0.05s
Gazelle-500@5%	14	0.01s	0.03s
Gazelle-500@1%	4405	0.20s	0.15s
Gazelle-500@0.5%	373568	$>$ 1000s	26.54s
JMLR-500@10%	636	0.04	0.08s
JMLR-500@5%	4751	0.04	0.17s
JMLR-500@1%	623011	0.09	7.39s
JMLR-500@0.5%	2782799	0.61	51.76s

TABLE II
NUMBER OF CLOSED PATTERNS AND RUNTIME ON EXISTING ALGORITHM FOR *Gazelle* AND *JMLR*.

B. Enumerating all closed frequent sequences

Specialized algorithms such as LCMseq [11] or Bide [22] have been designed to efficiently enumerate all closed frequent sequences. Run times for these algorithms are presented in Table II.

Our SAT-based approach does not purport to be faster than algorithms designed and optimized specifically for this basic task. Even then, it is worth noting that the number of closed frequent sequences grows quickly and makes the manual interpretation of the results almost impossible. (See Table II, second column). In this case, reducing the number of results by adding constraints is the only option offered to the user.

In some cases, post-processing the set of closed frequent sequences may be an alternative to using our constraint-based sequence mining framework. However, this option is not always available, and, when available, not always efficient.

For example, post-processing the pattern set in order to eliminate all the sequences that do not satisfy a $max\text{-}gap$ constraint is non-trivial: one needs to compute all the embeddings in the input dataset, and then check whether there exists one embedding satisfying the $max\text{-}gap$ constraint. Since the number of embeddings for a pattern is combinatorial with the number of valid matches for each character, there is no guarantee that this can be done in a reasonable amount of time.

Furthermore, the set of closed frequent sequences is not guaranteed to be a superset of closed frequent sequences satisfying a given constraint. As pointed by [19], this is true only for anti-monotonic constraints. In short, a frequent pattern that is not closed may become closed if a constraint discards its closed superset from the solution set. In this case, post-processing is simply impossible and existing sequence algorithms cannot be used.

C. Solving sequence mining with extra user constraints

Next, we demonstrate that, unlike specialized algorithms, the SAT solver is able take advantage of extra user constraints to reduce mining time. To do so, we compare run times based on various constraints:

- $max\text{-}gap$ and $max\text{-}span$ constraints, which regulate the embedding of sequences in the transactions;
- $regular$ constraints: for example, in the *JMLR* dataset, we extract the sequences matching the regular expression: "$\star machine \star learning \star$"[1];
- a combination of the above constraints.

The total run times are presented in Figure 3 for *JMLR-500* and Figure 5 for *Gazelle-500*.

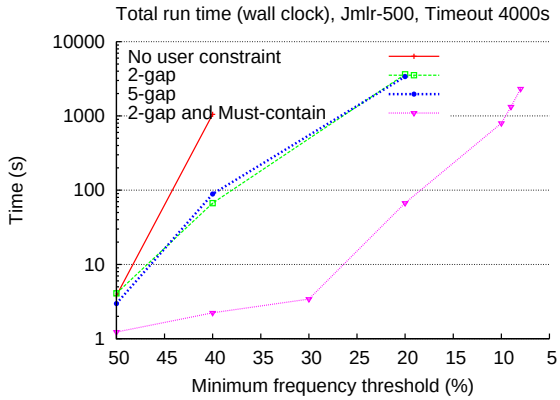

Fig. 3. Impact of user constraints and total solving times on *JMLR-500*.

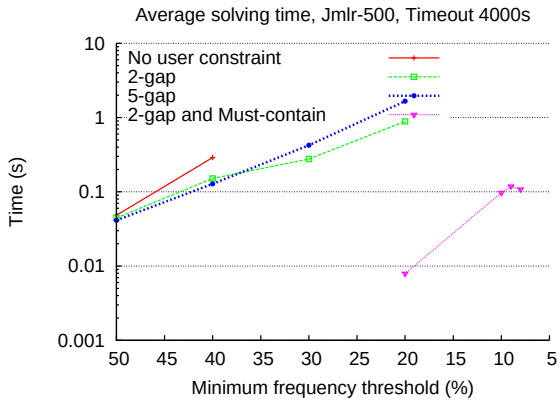

Fig. 4. Impact of user constraints and average solving times on *JMLR-500*.

It is important to note that, from the solver's point of view, adding constraints harden the problem but also reduces the number of required solver calls (Algorithm 1, line 3). In order to demonstrate that each individual call to the solver performs faster, we also measure the average solving time per run. The average solving times are presented in Figure 4 for *JMLR-500* and Figure 6 for *Gazelle-500*.

We first observe that, on *JMLR-500*, the gap constraint is very effective at reducing the total run time and the average solving time. However, on *Gazelle-500*, the gap and the span constraints have low impact. Indeed, the vast majority of sequences in *Gazelle-500* are very small in size (c.f. Table I). As a consequence, gap and span constraints do not reduce the number of solutions. Clearly, however, the addition of extra constraints does not negatively impact performance.

[1] Note this very simple regular constraint only make sense when they are combined with other constraints such as gap constraints. Otherwise a simple pre-processing removing all the transactions not containing the characters would be sufficient.

The combination of the regular constraints with the 2-gap constraints is very effective, allowing the solver to drastically reduce mining times.

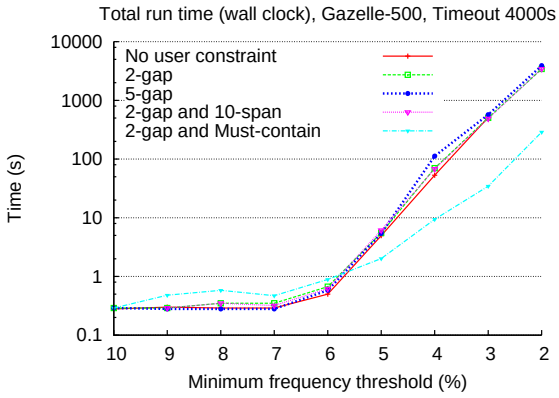

Fig. 5. Impact of user constraints and average solving times on *Gazelle-500*. Must-contain constraint is *must contain vocable-16*.

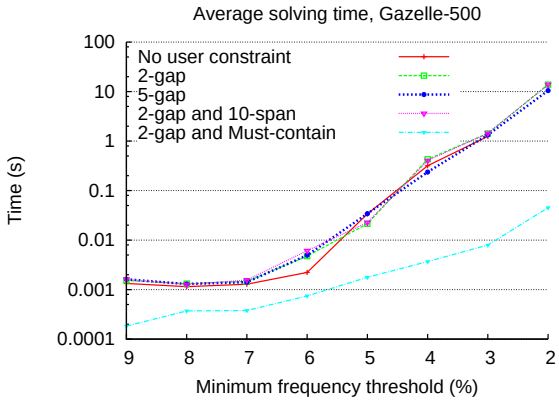

Fig. 6. Impact of user constraints and average solving times on *Gazelle-500*. Must-contain constraint is *must contain vocable-16*.

VII. Related work

Research on sequence mining was initiated by Agrawal et al.[23], who proposed an Apriori-derived algorithm to mine frequent sequences of itemsets (as opposed to mining frequent sequences of single characters, which is the issue that we address in this paper). The problem of mining sequences of single characters was introduced by Mannila et al. in [24] as the problem of *serial episode mining*. The authors proposed an Apriori-based algorithm with a depth-first search strategy. However, depth-first search *level wise* pattern mining algorithms usually require great memory capacity, because of the need to store all the patterns of size n before enumerating patterns of size $n+1$. An important improvement was therefore proposed by Wang et al. in [22] with the first depth-first search BIDE algorithm. In addition, these authors proposed a definition of closed patterns for sequence mining. The techniques introduced by BIDE enabled it to perform faster and to lower memory requirements by several orders of magnitude.

Attempts to improve the flexibility of existing sequence mining algorithms have been few. In [12], the authors mine sequences that satisfy a regular expression. The LCMseq algorithm ([11]) also deserves mention: It was extended by its authors to address different variations of the sequence mining task, such as maximal sequences and sequences with gap constraints. Unfortunatelly, it was proven incorrect (Section IV). Moreover, the above approaches for mining sequences with extra user constraints suffer from an important limitation: they can only support a small set of built-in constraints, and they fail to propose a general framework for the addition of extra constraints.

A more flexible approach was proposed by Guns et al. ([4]). In their work, these authors demonstrated that CP solvers and CP modeling techniques can be used to efficiently address various types of itemset mining problems. Since then, data mining and SAT/CP experts have obtained more results in the same vein. For example, Metivier et al. have proposed a constraint-based language for itemset mining ([25]). More recently, in [8], Coquerry et al. have proposed to perform simple sequence mining tasks using a SAT formulation.To the best of our knowledge, this represents the first attempt to address the sequence mining problem with the use of a SAT solver. However, this formulation is based on a set encoding of sequences proposed by Arimura and Uno ([26]) and cannot be extended to more complex mining tasks such as flexible sequences.

VIII. Conclusion & Perspectives

In this paper, we first proposed SAT encoding to adress the sequence mining problem. Second, we formulated several constraints, such as gap or span constraints, which are relevant to mining in various application fields. third, we proposed a sound methodology for addressing complex mining tasks, such as closed and maximal pattern mining. Further, we implemented this methodology with the Glucose SAT solver using interactive SAT solving.

To date, while state-of-the-art sequence mining algorithms are faster at generating all the frequent sequences, this approach is not scalable: since the number of sequences grows exponentially with the size of the dataset, unconstrained mining of sequences is bound to intractability. Our study and our experiments not only demonstrate the feasibility of using SAT for solving sequence mining problems, but also highlight several important benefits for data-mining users, including the ability to combine any arbitrary constraints.

In [27], the authors showed that highly optimized pattern mining algorithms are inherently difficult to parallelize due to the large number of instances of memory access. Using a SAT solver to tackle mining problems is a very different approach involving different algorithmic behavior and memory access patterns. Studying how our SAT-based approach may resolve the memory access problem presents an interesting research challenge.

REFERENCES

[1] R. Agrawal, R. Srikant *et al.*, "Fast algorithms for mining association rules in large database," in *VLDB*, vol. 1215, 1994, pp. 487–499.

[2] K. Ye, W. A. Kosters, and A. P. IJzerman, "An efficient, versatile and scalable pattern growth approach to mine frequent patterns in unaligned protein sequences," *Bioinformatics*, vol. 23, no. 6, pp. 687–693, 2007.

[3] N. Tatti and J. Vreeken, "The long and the short of it: summarising event sequences with serial episodes," in *KDD*, 2012, pp. 462–470.

[4] L. D. Raedt, T. Guns, and S. Nijssen, "Constraint programming for itemset mining," in *KDD*, 2008, pp. 204–212.

[5] O. Bailleux and Y. Boufkhad, "Efficient cnf encoding of boolean cardinality constraints," in *CP*, 2003, pp. 108–122.

[6] C. Sinz, "Towards an optimal cnf encoding of boolean cardinality constraints," in *CP*, 2005, pp. 827–831.

[7] I. Abío and P. J. Stuckey, "Conflict directed lazy decomposition," in *CP*, 2012, pp. 70–85.

[8] E. Coquery, S. Jabbour, L. Saïs, and Y. Salhi, "A sat-based approach for discovering frequent, closed and maximal patterns in a sequence," in *ECAI*, 2012, pp. 258–263.

[9] G. Audemard and L. Simon, "Predicting learnt clauses quality in modern sat solvers," in *IJCAI*, 2009, pp. 399–404.

[10] E. Giunchiglia, M. Maratea, and A. Tacchella, "Dependent and independent variables in propositional satisfiability," in *JELIA*, 2002, pp. 296–307.

[11] H. Ohtani, T. Kida, T. Uno, H. Arimura, and H. Arimura, "Efficient serial episode mining with minimal occurrences." in *ICUIMC*, 2009, pp. 457–464.

[12] H. Albert-Lorincz and J.-F. Boulicaut, "A framework for frequent sequence mining under generalized regular expression constraints," in *Workshop on Knowledge Discovery in Inductive Databases (KDID-2003)*, 2003, p. 2.

[13] G. Pesant, "A regular language membership constraint for finite sequences of variables," in *CP*, 2004, pp. 482–495.

[14] C.-G. Quimper and T. Walsh, "Global grammar constraints," in *CP*, ser. Lecture Notes in Computer Science, F. Benhamou, Ed., vol. 4204. Springer, 2006, pp. 751–755.

[15] M. Sellmann, "The theory of grammar constraints," in *CP*, ser. Lecture Notes in Computer Science, F. Benhamou, Ed., vol. 4204. Springer, 2006, pp. 530–544.

[16] C.-G. Quimper and T. Walsh, "Decompositions of grammar constraints," in *AAAI*, D. Fox and C. P. Gomes, Eds. AAAI Press, 2008, pp. 1567–1570.

[17] G. Katsirelos, N. Narodytska, and T. Walsh, "Reformulating global grammar constraints," in *CPAIOR*, 2009, pp. 132–147.

[18] N. Pasquier, Y. Bastide, R. Taouil, and L. Lakhal, "Discovering frequent closed itemsets for association rules," *ICDT99*, pp. 398–416, 1999.

[19] F. Bonchi and C. Lucchese, "On closed constrained frequent pattern mining," in *ICDM*. IEEE, 2004, pp. 35–42.

[20] R. Bayardo, "Efficiently mining long patterns from databases," in *ACM SIGMOD*, 1998, pp. 85–93.

[21] A. Nadel and V. Ryvchin, "Efficient sat solving under assumptions," in *SAT*, A. Cimatti and R. Sebastiani, Eds., 2012, pp. 242–255.

[22] J. Wang and J. Han, "Bide: Efficient mining of frequent closed sequences," in *ICDE*. IEEE, 2004, pp. 79–90.

[23] R. Agrawal and R. Srikant, "Mining sequential patterns," in *ICDE*. IEEE, 1995, pp. 3–14.

[24] H. Mannila, H. Toivonen, and A. Inkeri Verkamo, "Discovery of frequent episodes in event sequences," *DMKD*, vol. 1, no. 3, pp. 259–289, 1997.

[25] J. Metivier, P. Boizumault, B. Crémilleux, M. Khiari, and S. Loudni, "A constraint-based language for declarative pattern discovery," in *Data Mining Workshops (ICDMW)*. IEEE, 2011, pp. 1112–1119.

[26] H. Arimura and T. Uno, "Polynomial-delay and polynomial-space algorithms for mining closed sequences, graphs, and pictures in accessible set systems," in *ICDM*, 2009, pp. 1087–1098.

[27] B. Negrevergne, A. Termier, M.-C. Rousset, and J.-F. Méhaut, "Para miner: a generic pattern mining algorithm for multi-core architectures," *Data Mining and Knowledge Discovery*, pp. 1–41, 2013.

[28] F. Benhamou, Ed., *Principles and Practice of Constraint Programming - CP 2006, 12th International Conference, CP 2006, Nantes, France, September 25-29, 2006, Proceedings*, ser. Lecture Notes in Computer Science, vol. 4204. Springer, 2006.